

AMENDMENTS TO THE SPECIFICATION

Please replace the paragraph starting at line 3, page 2 and ending at line 8, page 2, with the following amended paragraph:

Broadband Internet network infrastructure is developing at rates that exceed even aggressive ~~analyst's~~ analysts' predictions. In the consumer market sector, telecommunications, cable and wireless companies have accelerated deployment of broadband capability to the home with xDSL, cable modem or wireless last mile rollouts. In the corporate market sector, broadband infrastructure is already available for desktop computing applications.

Please replace the paragraph starting at line 5, page 3, and ending at line 20, page 3, with the following amended paragraph:

As Broadband connections proliferate, demand for better performance has fostered an industry focused on speeding up the delivery of Internet content. The majority of these solutions have centered on smaller objects such as text and images. Video data or objects present problems due to their data size and the requirement to provide it them at a particular rate related to the real-time or near-real time play or rendering requirement. Due to its sheer size alone, video is one of the most difficult data types to manage on the Internet or other network environment. A five-minute video clip, encoded and compressed at 1.5 Mbps is 56 Megabytes in size. This compares to the few kilobyte data sizes for typical web pages. The strict video timing requirements impose additional constraints. When a frame or set of frames arrive past their intended presentation time (for example, at greater than the nominal 1/30 second frame interval in the case of a 30 fps video) the consumer or user experiences jerky playback, dropped frames or segments of the video, or other defects that detract from the viewability of the video and render it essentially useless in a commercial setting. Given these stringent requirements, delivering quality video over broadband is a challenging problem.

Please replace the paragraph starting at line 9, page 4, and ending at line 14, page 4, with the following amended paragraph:

Existing conventional solutions geared towards improving the performance of accessing web pages containing rich media (typically including static images) are increasingly being used to address the problems with streaming video on the Internet. Currently, there are two classes of solutions that have been employed for improving performance of content distribution on the Internet: (i) particular content delivery network architectures and operational schemes, and (ii) ~~and~~ content caching schemes.

Please replace the paragraph starting at line 16, page 5, and ending at line 6, page 6, with the following amended paragraph:

When the content is video, a third major problem or limitation with such direct delivery becomes evident. Contemporary networks are packet switched and there may typically be a number of routers and switches between the central object server and the requesting user. Routers are typically provided with buffers for buffering data (usually in the form of packets) received until it can be forwarded to the next node in the network, however these buffers have limited buffering capacity and in the event that the amount of data received is greater than the amount that can be forwarded or stored until forwarding is possible, such data or packets of data may simply be lost or dropped. For typical web pages, this does not represent a severe problem as the page is ~~simple~~ simply requested again. However, for a video stream intended to be viewed continuously and in real time, dropping of a packet of video does not provide any recovery mechanism. That segment of video simply cannot be viewed and various schemes may be provided to substitute for that video segment, such as a static freeze of the last available frame, blanking the screen, or other conventional but usually unsatisfactory techniques.

Please replace the paragraph starting at line 6, page 11, and ending at line 15, page 11, with the following amended paragraph:

One of the benefits of the Internet or computer networks is its ability to provide "narrowcasting" – for example, ability to address small groups of users (and single users) in a targeted manner. The promise of narrowcasting is in its ability to provide targeted information to an end user. In the broadcast world (for example, network television), all users tuned to a particular program (for example, the NBA finals) ~~receives~~ receive the same program, including

the same advertisements. In a narrowcasting world, it should be possible for a user in Cincinnati, interested in automobiles to be seeing advertisements from car dealerships in their local area. This would mean that the information about the user ("metadata" or "MD") 108 be available at the edge server 110 for dynamic content insertion.

Please replace the paragraph starting at line 18, page 12, and ending at line 1, page 13, with the following amended paragraph:

With respect to the delivery of personalized information on demand, personalized or customized delivery of information rich in video content (~~new news~~, sports, entertainment, personal health information, and other types of video-rich content-) is a growing application segment on the broadband Internet. A five-minute video segment at 1.0 Mbps amounts to 37.5 megabytes. One such channel of video, which is a 24-hour segment split into 5 minute segments amount to about 10 gigabytes of storage. A hundred such channels amount to 1 terabyte. Such media stored on 1000 edge servers amount to 1 petabyte of storage for one day's worth of video.

Please replace the paragraph starting at line 6, page 15, and ending at line 15, page 15, with the following amended paragraph:

Method, system, computer program and computer program produce for a metadata enabled push-pull model and method for efficient low-latency video-content distribution over a network are provided. Metadata is used as a vehicle and mechanism to enable intelligent decisions to be made on content distribution system operation. Metadata is data that contains information about the actual content, and in some cases, the metadata may also contain portions of the content or a low-resolution preview of the content. Aspects of the invention are directed toward the distribution of metadata throughout the network in a way that facilitates efficient system operation as well as optionally but advantageously providing a set of services such as tracking, reporting, personalization, and the like.

Please replace the paragraph starting at line 16, page 15, and ending at line 26, page 15, with the following amended paragraph:

In one embodiment, the invention provides a metadata enabled server for distributing a content object to a user over a network in response to a user request, the metadata enabled server including: a computer having a processor and a memory coupled to the processor for executing computer program instructions, and at least one input/output port for receiving and sending communications from external entities; a storage device coupled to the server and storing metadata describing content objects accessible to the server including a location from where a content object is stored and may be directed to the user; and a controller for distributing the content object to the user using the metadata and maintaining isochronous delivery of portions of the content over the network. Method and procedures, system, and computer program for distributing content and controlling distribution of content are also provided.

Please replace the paragraph starting at line 20, page 16, and ending at line 7, page 17, with the following amended paragraph:

The basic tenet of at least some embodiments of the invention described herein is the use of metadata as a vehicle and mechanism that enables intelligent decisions to be made on system operation. Metadata contains information about the actual content: for example, its physical properties, possible locations of the content represented by the metadata, its usage terms, and the like, and others as described in greater detail elsewhere in this description. In some cases, the metadata may also contain portions of the content ("content prefix") or a low-resolution preview of the content. Aspects of the invention are directed toward the distribution of metadata throughout the network in a way that facilitates efficient system operation as well as optionally but advantageously providing a set of services (tracking, reporting, personalization, and the like) that are not present in the conventional or prior-art systems and methods.

Please replace the paragraph starting at line 8, page 17, and ending at line 6, page 18, with the following amended paragraph:

With reference to the illustration in FIG. 3, there is shown in block diagram form ~~of~~ an embodiment of ~~the~~ system 102 of the present invention as deployed in a global computer network. ~~An origin~~ Origin server 104 may be a server of which many types and configurations are known in the art. Typically such servers include a processor, memory coupled to the

processor, input/output device, and mass non-volatile storage usually in the form of rotating media hard disk drives, and computer software providing computer executable instructions that execute in the processor and memory to direct the server to operate in a particular manner. In this embodiment, ~~the~~ origin server 104 includes a metadata database 106 and a video content (or other arbitrary content) store 108. System 102 also includes at least one and typically a plurality of edge servers (ES) 110-n (e.g. 110-1, 110-2, ..., 110-n) each of which also typically includes a metadata database 106 and mass storage 118-n (e.g., 118-1, 118-2, ..., 118-n). The data stored on each edge server mass storage device 118-n may usually depend on the particular implementation as described in greater detail herein elsewhere. For example, in one embodiment ~~the~~ edge server storage 118-n will store full video content once at least a single request has been made for the particular content. In other embodiments, ~~the~~ edge server storage 118-n will store prefix content such as a video prefix content portion. In other embodiments, ~~the~~ edge server storage 118-n may store both full video content or substantially full video content for a plurality of video titles as well as a prefix portion for some set of video content items. ~~The metadata~~ Metadata database 106 and the content may be stored on the same physical and/or logical device or on different physical or logical mass storage devices. As ~~the~~ metadata database 106 is relatively small, it may instead or in addition be stored in random access memory, typically with non-volatile backup storage.

Please replace the paragraph starting at line 7, page 18, and ending at line 15, page 18, with the following amended paragraph:

With respect to FIG. 4, there is shown in block diagram form an embodiment of a media store including media streaming support 122 for MPEG-1 and MPEG-2 pumps, Kasenna (Mountain View, California) OT 4.0 format pump, real G2 server, and support for other media service to support a variety of video ~~object~~ objects and other content types and formats. Media management 124 functionality is provided, including functionality for acquisition, storage, and metadata database and data management. Media distribution 126 functionality is also provided, including scheduled transfers, on-demand transfers, and uni-cast and multi-cast operation. Connection manager 116 and storage manager 117 functionality may also be provided.

Please replace the paragraph starting at line 16, page 18, and ending at line 20, page 18, with the following amended paragraph.

With respect to FIG. 5, there is shown in block diagram form of the major components of a computer ~~such~~ as may be used in conjunction with the system of the present invention to receive and render content received. A communication connection provides a communication link or path for receiving content, such as a video content stream, from a server.

Please replace the paragraph starting at line 21, page 18, and ending at line 4, page 19, with the following amended paragraph:

By way of a top-level description, in one particular embodiment of the inventive system and method, edge servers or "ES" (that is servers near the edge of the network that are primarily responsible for serving content to end users) on the content distribution network maintain a directory that organizes metadata for individual content items typically in the form of a universal reference locator (URL) for the item. The directory may be part of a relational database in which functions are provided that allow easy manipulation of the metadata, or simpler flat file database structures or lists or ~~table~~ tables may be utilized.

Please replace the paragraph starting at line 5, page 19, and ending at line 11, page 19, with the following amended paragraph:

In one embodiment, video content URLs, which reference the location of the video objects or assets, regardless of the type of media format it ~~they represents~~ represent, are mapped in a homogeneous format. User selection of particular video content whether by URL or by other designator or identifier may be recorded in the metadata database, regardless of whether the video object is cached or not. Tracking, reporting, billing, targeted advertising, and other data collection and mining operations, are easily enabled using this approach.

Please replace the paragraph starting at line 13, page 20, and ending at line 5, page 21, with the following amended paragraph:

In another aspect, video prefix caching or video prefix content distribution (VPCD) may be employed in a manner that involves distributing metadata and beginning portions of a video

object to edge servers based on characteristics such as anticipated demand, measured usage, and type of connection between the origin server (OS) 104 and edge server (ES) 110. For example, large video prefixes can be used for edge servers in which greater demand is anticipated, and smaller prefixes stored on other edge servers. If a hit occurs for a video object having a prefix on the edge server, streaming can begin upon demand, while the remaining portion of the video object is fetched and cached. Typically, the origin server 104 will provide a video content store (VCS) 108 or other type of content store, storing the video or other content for all of the content items. Each edge server 110 may also provide a video or other content store for storing content items, though each edge server's content store will typically store only a subset of the content items stored in the origin server's content store. (It may be noted that certain embodiments of the invention described hereinafter in the context of virtualization may not require a content store ~~118~~ 108 or will require a much smaller one used, that may for example, be used to buffer the incoming video prior to transmitting it to the requesting user.)

Please replace the paragraph starting at line 24, page 23, and ending at line 12, page 24, with the following amended paragraph:

The metadata may also define or specify the "cost" associated with providing the content from a first (source) location on a network to another second location on the network, the second location usually being the edge or other server that will serve the content to the requestor. Where separate networks are involved, the cost may include the cost of crossing over from one network to the other. Cost factors, may for example include one or more of tariffs for use of network bandwidth, cost of storage on a particular locations of data storage devices, requirements to use external entities or avoid use of external entities, as well as other factors associated with transport and/or storage. Because the cost associated with content retrieval and transport from location to location may change according to network conditions or other factors, an external agent having access to network conditions, network tariffs, contractual information, current bidding forms, or other facts that would influence cost and suggest alternative locations and/or network routings may be employed.

Please replace the paragraph starting at line 7, page 25, and ending at line 24, page 25, with the following amended paragraph:

In at least one embodiment the metadata consists of (i) a globally unique identifier that identifies and optionally describes the content associated with the metadata, and (ii) a location or locations at which the content may be found, usually in the form of at least the origin server identifier and more typically also the identities of edge servers that cache the content. Other ~~of the~~ metadata elements may optionally but advantageously be provided to support other desirable content distribution features. By way of highlight and summary, these other optional metadata elements may be selected in any combination from the set: (iii) cost of retrieving the content from each location; (iv) content format such as MPEG version, RealVideo, or other ~~known~~ known or to be developed audio and/or video formats; (v) rights information; (vi) time-to-live or expiration information; (vii) quality-of-service; (viii) content subset of preview derivative; such as single frame video, low-resolution video, limited length video, textual description, or the like; and/or (ix) content prefix such as defined time duration of actual full video. The invention further contemplates that the metadata may include any other data or information that describes the content, assists in its localization or routing to a requestor, controls access to the content, assists in maintaining a desired level of quality, or otherwise reduces per server or total storage requirements, or reduces latency.

Please replace the paragraph starting at line 12, page 26, and ending at line 17, page 26, with the following amended paragraph:

Metadata is distributed throughout the network and portions of the metadata are stored in servers that receive metadata updates. Since typical metadata for any particular video are orders of magnitude smaller than the video data files themselves, distribution of metadata in the network and storage at the edge servers (or even at clients directly) is a viable operation and does not significantly impact edge servers storage.

Please replace the paragraph starting at line 6, page 27, and ending at line 11, page 27, with the following amended paragraph:

Attention is now directed to an embodiment of a procedure for installing content data and metadata associated with the content data on the system, network, and/or components thereof. This procedure is described relative to the diagrammatic flow chart illustration in FIG. 6. First, it

is necessary to install, generate, or otherwise place the content (video) into the system (~~Step 201~~) and if not initially installed onto the origin server then install onto the origin server using known techniques.

Please replace the paragraph starting at line 12, page 27, and ending at line 20, page 27, with the following amended paragraph:

Next, create the metadata for the content (~~Step 202~~). This metadata creation may be done at any time including before, during, or after the content (video) is created so long as the content is sufficiently well defined to permit generation of the metadata for the content. The metadata at this stage may include only the globally unique ID and an origin server location identifier. The origin server location is identified in the metadata even though the video may not actually have been installed on the origin server at that time because it will have been installed there by the time the metadata is actually used or queried. The metadata may subsequently be modified or updated to reflect changes in content storage location.

Please replace the paragraph starting at line 21, page 27, and ending at line 7, page 28, with the following paragraph:

Supplemental metadata elements may optionally be added as required or desired to provide optional features, capabilities, and performance (~~Step 204~~). For example, the metadata may be augmented with other than a minimum set of elements to identify rights, format, or other of the characteristics of the content as described herein above. This additional metadata may optionally be packaged with or attached to the content and extracted from the content thereby eliminating any manual or separate upload/download steps. The metadata may alternatively be manually entered such as from a text form, uploaded or downloaded from an external source, or automatically be extracted from some file, such as the content data file or any other file. For example, for video content in the MPEG-7 format, metadata may be included with the file and extracted from it.

Please replace the paragraph starting at line 8, page 28, and ending at line 19, page 28, with the following amended paragraph:

A push, and desirably a “scheduled push”, of the metadata to the edge servers is performed (~~Step 206~~). In one embodiment, the metadata is pushed to all edge servers. In another embodiment, the metadata is pushed to a selected subset of edge servers. Note that the full content (typically Megabytes or Gigabytes) is not being pushed to the edge servers, only the metadata (typically hundreds or a few thousands of bytes) is being pushed and represents orders of magnitude less data than the content itself for most video content of interest here. A scheduled push is a push performed according to a specific schedule or schedule policies and is usually designed to minimize cost and/or disruption to other network activity. Conveniently it may be performed during non-peak hours (overnight) and coordinated in time over the system so that the push is not attempted to all edge servers simultaneously but spread out over time to reduce peak bandwidth needs and server or other network node loading.

Please replace the paragraph starting at line 4, page 29, and ending at line 11, page 29, with the following amended paragraph:

In another optional embodiment, when the system and method provides or utilizes one or more servers other than the origin server and the edge server, such as one or more directory servers, the content metadata may also be pushed to the directory server (~~Step 210~~). This configuration is optionally provided as an optimization of the basic system and method and provides a single location data structure that may be queried to identify all locations at which a content item is located. Directory servers may be replicated at a plurality of locations throughout the network.

Please replace the paragraph starting at line 12, page 29, and ending at line 15, page 29, with the following amended paragraph:

The metadata is installed on the system. (It is noted that metadata may be installed and/or updated on the system without reinstalling the video or other content.) Once the scheduled push has been performed, the metadata is updated to reflect the additional content and metadata changes (~~Step 212~~).

Please replace the paragraph starting at line 21, page 29, and ending at line 4, page 31, with the following amended paragraph:

It is noted that this hybrid “push” and “pull” with metadata based method and system have many advantages over conventional systems and techniques. In a conventional pull method, there is no metadata and no tracking capability is available as to content. In the system described here, metadata data records contain fields to track a variety of characteristics and parameters, for example embodiments of the invention may provide for ~~meta-data~~ metadata records containing fields to track one or a combination of: the number of users that accessed the content, the number of users that were denied access to the content, the frequency of access, the length of viewing, or the like. Such information can then be used for data mining purposes: for example, if the frequency of access of a particular content goes over a preset threshold, the system can trigger an operation that eventually results in the “push” of the content from an origin server (or any other server) to the rest of the edge servers (or to some sub-set of the rest of the edge servers). The collected metadata is sent back to an origin server or any other server, either in real-time or packaged and sent periodically or according to any other desired schedule. This optional tracking capability stores user access data (possibly individual user data but more typically statistical data in the interests of maintaining user privacy) in metadata records on the edge server, or in any external location designated for gathering, collection, and/or analysis of the user access data. Typically, the user access data would be gathered at each of the edge servers, and might include by way of example but not limitation, the number of requestor’s for each content item, the time of day that the request was made, the delay period between the request and the satisfaction of the request with playback, the sufficiency or excess of context prefix information, the frequency of user disconnects prior to the content being delivered, ratings information or information from which ratings information may be derived, frequently requested content items, items in the directory that are infrequently requested, and any other information that may assist in tuning the network configuration and content delivery, selecting content and content formats, or otherwise tuning or optimizing performance. This collected user access and performance data (a type of metadata in itself) is advantageously stored and pushed back to a processing location for analysis and reporting. The analysis may include various data mining techniques, statistical analysis, or other tuning and/or optimization techniques.

Please replace the paragraph starting at line 8, page 32, and ending at line 21, page 32, with the following amended paragraph:

For networks having sufficiently high bandwidth and availability communication channel links or pipes between the origin server and the edge server, a further ~~embodiments~~ embodiment eliminates or substantially eliminates local content storage or caching of complete content items (e.g. full videos) on the edge servers. Rather, once a user request is made to an edge server, the edge server pulls the requested content from the origin server and initiates playback to the requesting user once a sufficient quantity of the requested content has been received. The network is configured in such a manner that resources are available for such real-time delivery with some predetermined or specified probability. Optionally, but desirably a initial or prefix portion of each content item is cached so that playback to the requesting user can begin with little or no delay. The length or playback duration of the prefix may be adjusted according to measured and/or predicted delay associated with a content request by a particular edge server to a content source such as the origin server or other content servers or cache servers that may be distributed in the network.

Please replace the paragraph starting at line 1, page 34, and ending at line 9, page 34 with the following amended paragraph:

An embodiment of the request response and playback procedure 300 is now described. A user submits an implicit or explicit request (Step 301) for a content item (such as for a video V) to a web site, portal, or other network access point. This user is referred to as the user, user requestor, or simply as the requestor. An entity within the network receives the request and if this receiving entity is not an edge server, forwards it to an edge server. We assume for purposes of description that the edge server receives the request directly from the user and ~~ignore~~ ignores any other intermediaries, service providers, routers, or the like that may actually be interposed between the requestor and the edge server.

Please replace the paragraph starting at line 10, page 34, and ending at line 8, page 35, with the following amended paragraph:

The edge server receives the user request for video V (Step 302) and determines whether its local directory includes a directory entry for video V (Step 304). In one embodiment, the local directory stores entries -indicating the availability of the content, the rights to the content,

and the like. When a user selects a video title displayed on his computer, TV set via a set-top box, or other information appliance or access method, the user essentially transmits a URL or other globally unique identification for the video (or other content item) to the network. Intelligence embodied in the network (such as proximity of the user to an edge server near his/her geographic area) directs the request to that server. The edge server receiving the request examines its directory to determine if the metadata for the video V is in ~~it's~~ its directory (Step 304). If the metadata is located, it will indicate whether the video V is completely stored on the edge server storage, whether video V is currently cached inside the computer's main memory or other storage, or whether video V is only partially available (for example, the video prefix only). If only the metadata for video V is available, it will indicate where in the network copies of video V may be available (either at video V's origin server or on other edge servers that have fetched video V and is are willing to serve it to our particular edge server). If it is determined that the metadata for video V is not in the receiving edge server's directory, an optional query is made to an external directory or directories for the selected video V (Step 312). These external directories may for example be selected from one or more of edge servers themselves or may be independent directories created for the purpose of making locating the videos easier (such as the Domain Name Service implementing Internet host name resolutions), among others.

Please replace the paragraph starting at line 14, page 37, and ending at line 6, page 38, with the following amended paragraph:

Once a source location (and optionally an "optimum" location given the set of possible locations) for the video V has been identified (Step 320) a determination as to whether network resources are available to deliver the video from the identified source server and the requesting edge server is made (Step 322). These resources are then reserved (Step 326). If it is determined that adequate resources are not available and cannot be reserved within a predetermined (or otherwise determined) acceptable period of time, either the request fails (Step 324) and notification of such failure is communicated to the requestor or preferably, an additional determination is made to identify a new source (Step 320). In practice, the determination as to whether resources are available (Step 322) may be part of the step of finding a source location (Step 320) or a separate step. Advantageously, it is combined so that identifying the best or

optimal source location takes into account the available network resources, and the quality of service of the available network resources. In other embodiments, performing this as a separate or even as an additional check may be desirable where network resource availability may change rapidly so that the check would be made just prior to initiating a pull of the video V from the identified source location back to the requesting edge server (Step 326 328).

Please replace the paragraph starting at line 15, page 38, and ending at line 4, page 39, with the following amended paragraph:

Once the source location has been identified and network resource availability verified, the edge server begins pulling the video V content from the identified source destination (Step 328). Because it will take some period of time for all of the video to be communicated from the source to the edge server, once the process has begun, a determination is made as to whether there is a sufficient amount of the requested video V so that playback to the requesting user may be initiated and maintained according ~~in~~ to an acceptable manner. The acceptability criteria may be established in a variety of ways, but for example may be set so that there are no discontinuities in playback, maintenance of predetermined playback rate, providing set QOS, or otherwise establish to maintain desired video quality. In general, the amount of video required to have been pulled may depend on the nature, quality, and availability of the network resources between the sending source server and receiving edge server. Typically, a greater amount of received video ~~being~~ is required for poorer intervening network resources than for better intervening network resources.

Please replace the paragraph starting at line 23, page 39, and ending at line 8, page 40, with the following amended paragraph:

Once the full video V has been received by the edge server, the metadata within the edge server is updated to identify that the full video V is now cached within the edge server (Step 332). Optionally, but desirably, the metadata in other of the edge servers is updated either at the same time, or as part of a scheduled update (Step 324). The changed ~~meta-data~~ metadata is also desirably communicated to the origin server for storage there, and to the extent that the origin server is responsible for administering metadata throughout the system, for subsequent

Application No: 10/090,709

Reply to office action mailed on March 24, 2005

Art Unit: 2142

communication to and storage by the other edge servers. In embodiments of the invention that provide a directory server, the modified metadata (or an indication of the particular change in the metadata) is communicated to and stored in the directory server.